

Evolution of cooperation facilitated by reinforcement learning with adaptive aspiration levels

Shoma Tanabe¹ and Naoki Masuda^{2,3,*}

¹ Faculty of Engineering, University of Tokyo
7-3-1 Hongo, Bunkyo, Tokyo 113-8656, Japan

² Graduate School of Information Science and Technology
University of Tokyo, 7-3-1 Hongo, Bunkyo, Tokyo 113-8656, Japan

³ PRESTO, Japan Science and Technology Agency
4-1-8 Honcho, Kawaguchi, Saitama 332-0012, Japan

* Corresponding author (masuda@mist.i.u-tokyo.ac.jp)

November 8, 2011

Abstract

Repeated interaction between individuals is the main mechanism for maintaining cooperation in social dilemma situations. Variants of tit-for-tat (repeating the previous action of the opponent) and the win-stay lose-shift strategy are known as strong competitors in iterated social dilemma games. On the other hand, real repeated interaction generally allows plasticity (i.e., learning) of individuals based on the experience of the past. Although plasticity is relevant to various biological phenomena, its role in repeated social dilemma games is relatively unexplored. In particular, if experience-based learning plays a key role in promotion and maintenance of cooperation, learners should evolve in the contest with nonlearners under selection pressure. By modeling players using a simple reinforcement learning model, we numerically show that learning enables the evolution of cooperation. We also show that numerically estimated adaptive dynamics appositely predict the outcome of evolutionary simulations. The analysis of the adaptive dynamics enables us to capture the obtained results as an affirmative example of the Baldwin effect, where learning accelerates the evolution to optimality.

1 Introduction

The mechanisms of cooperation in social dilemma situations are a central topic in interdisciplinary research fields including evolutionary biology, ecology, economics, and sociology. As analyzed by the prisoner's dilemma (PD) game and its relatives, direct reciprocity is among the main known mechanisms underlying cooperative behavior (Trivers, 1971; Axelrod, 1984). In direct reciprocity, iterated interaction between the same individuals motivates them to continue cooperating (C) rather than to defecting (D) to obtain momentarily large payoffs; defection would be negatively rewarded by the opponent player's retaliation in later rounds. Variants of the celebrated retaliatory strategy tit-for-tat (mimicking the opponent's action in the previous round) (Nowak & Sigmund, 1992) and a win-stay lose-shift strategy (Kraines & Kraines, 1989; Nowak & Sigmund, 1993) are recognized as strong competitors in the iterated PD game.

In the iterated games concerning direct reciprocity, it is natural to assume that players modify their strategies in response to their experiences in past rounds. The tit-for-tat, its variants, and win-stay lose-shift strategies can be interpreted as examples of such learning strategies because the tit-for-tat, for example, implies that the player selects the action (i.e., C or D) depending on the result of the last round. A more sophisticated learning player of this kind exploits a longer history of the game for action selection (e.g., cooperate if the player and the opponent cooperated in the previous two rounds, and defect otherwise) (Lindgren, 1991). Classes of other learning models include fictitious play and reinforcement learning (Camerer, 2003; Fudenberg & Levine, 1998). Learning apparently seems beneficial in iterated games because learning players are more flexible than nonlearning players.

If learning is a key factor in promoting cooperation in real societies, the number of learning players should increase when a population evolves under selection pressure. However, the advantage of learners over nonlearners in evolutionary dynamics is elusive because a pair of learning players often results in mutual defection (Macy, 1996; Sandholm & Crites, 1996; Posch *et al.*, 1999; Taiji & Ikegami, 1999; Macy & Flache, 2002; Masuda & Ohtsuki, 2009) and learning may be costly.

The constructive roles of learning in the evolution of certain traits are collectively called the

Baldwin effect (see Simpson, 1953; Turney *et al.*, 1996; Weber & D. J. Depew, 2003; Crispo, 2007; Badyaev, 2009 for reviews). Although earlier examples of the Baldwin effect are not necessarily founded on firm empirical evidence (Simpson, 1953; Weber & Depew, 2003), there exists a plethora of positive evidence of the Baldwin effect. Examples include fly's morphological developments (Waddington, 1942), colonization of house finch in North America (Badyaev, 2009), and persistence of coastal juncos (Yeh & Price, 2004). In fact, the concept of the Baldwin effect differs by authors (see Simpson, 1953; Downes, 2003; Turney *et al.*, 1996). Although earlier computational models suggest that learning accelerates evolution (Hinton & Nowlan, 1987; Ancel, 1999; Maynard Smith, 1987), later theoretical and numerical studies suggest that learning either accelerates or decelerates evolution toward the optimum depending on the details of the models (Ancel, 2000; Dopazo *et al.*, 2001; Borenstein *et al.*, 2006; Paenke *et al.*, 2007; Paenke *et al.*, 2009). The advantage of learning in evolution is also nontrivial in this broader context.

We numerically investigate the effect of learning on evolution in the iterated PD game. This question was explored in previous literature (Suzuki & Arita, 2004; also see Wang *et al.*, 2008 for discussion). Our emphasis in this study is to use a reinforcement learning model for the iterated PD game (Masuda & Nakamura, 2011) that is much simpler in terms of the number of plastic elements than the plastic look-up-table model adopted in (Suzuki & Arita, 2004). In our model, players are satisfied with and persist in the current action when the obtained payoff is larger than a plastic threshold. Our model of players introduced in (Masuda & Nakamura, 2011) modifies those in (Karandikar *et al.*, 1998; Posch *et al.*, 1999; Macy & Flache, 2002). Via the stability analysis for nonlearning players, the numerical analysis of the discretized adaptive dynamics with nonlearning and learning players, and full evolutionary simulations, we show that learning is needed for a noncooperative population to evolve to be able to engage in mutual cooperation for wide parameter ranges. We also discuss our results in the context of the Baldwin effect.

2 Model

2.1 Iterated PD game

We assume that each player plays the PD game against each of the other players in a population. In each round t ($t = 1, 2, \dots$) within a generation, a player selects C or D without knowing the action (i.e., C or D) of the opponent player. The payoff to the focal player is defined by

$$\begin{array}{cc} \text{C} & \text{D} \\ \text{C} & \begin{pmatrix} R & S \\ T & P \end{pmatrix} \\ \text{D} & \end{array}, \quad (1)$$

where $T > R > P > S$ and $R > (T + S)/2$. Equation (1) represents the row player's payoff. The payoff to the opponent (column player) is defined likewise; the PD game is symmetric. Because $T > R$ and $P > S$, mutual defection is the only Nash equilibrium of the single-shot PD game.

However, players may continue mutual cooperation for their own benefits in the iterated PD game (Trivers, 1971; Axelrod, 1984). We denote the number of rounds per generation by t_{\max} . Technically, the Nash equilibrium of the iterated PD game is perpetual mutual defection if the players know t_{\max} beforehand. The number of rounds is often randomized to avoid this effect (Axelrod, 1984). To simplify the analysis, we assume that the players are unaware of the fixed value of t_{\max} .

Earlier studies identified tit-for-tat, which involves imitating the previous action of the opponent, as a strong strategy in the iterated PD game when various strategies coexist in a population (Trivers, 1971; Axelrod, 1984). However, later studies showed that tit-for-tat is not robust against error and that alternative strategies such as generous tit-for-tat (Nowak & Sigmund, 1992) and Pavlov (Kraines & Kraines, 1989; Nowak & Sigmund, 1993) are strong competitors in the iterated PD game with error. By definition, a Pavlov player receiving payoff T or R is satisfied and does not change the action in the next round, whereas the same player receiving payoff P or S is dissatisfied and flip the action. A population composed of Pavlov players, for example, realizes mutual cooperation such that a player gains approximately R per round.

2.2 Reinforcement learning

Intuitively, the ability to learn may seem to be an advantageous trait in the iterated PD game if the cost of learning is negligible. However, this is generally not the case. A pair of learning players often ends up with mutual defection unless a learning algorithm is carefully designed (Macy, 1996; Sandholm & Crites, 1996; Posch *et al.*, 1999; Taiji & Ikegami, 1999; Macy & Flache, 2002; Masuda & Ohtsuki, 2009). Learning requires trial and error, i.e., the exploration of unknown behavioral patterns as well as the exploitation of known advantageous behavioral patterns. Exploratory behavior of a learning player may look just random to opponents, and it is rational to defect against random-looking players.

To compromise the possibility of mutual cooperation, the simplicity of the learning algorithm, and the biological plausibility of the model as compared to some other learning algorithms, we use a variant of the Bush–Mosteller (BM) reinforcement learning model (Masuda & Nakamura, 2009). This model modifies the models in the previous literature (Karandikar *et al.*, 1998; Posch *et al.*, 1999; Macy & Flache, 2002) such that players learn to mutually cooperate for wide parameter ranges.

In round t , the cooperability of the learning player is given by the probability p_t . We update p_t using the results of the single-shot PD game as follows:

$$p_{t+1} = \begin{cases} p_t + (1 - p_t)s_t & \text{(action in round } t \text{ is C, and } s_t \geq 0), \\ p_t + p_t s_t & \text{(action in round } t \text{ is C, and } s_t < 0), \\ p_t - p_t s_t & \text{(action in round } t \text{ is D, and } s_t \geq 0), \\ p_t - (1 - p_t)s_t & \text{(action in round } t \text{ is D, and } s_t < 0), \end{cases} \quad (2)$$

where

$$s_t = \tanh[\beta(r_t - A_t)], \quad (3)$$

and $r_t \in \{R, T, S, P\}$ is the payoff to the player in round t . s_t stands for the degree of satisfaction in round t . When s_t is large, the player increases the probability of taking the current action in round $t + 1$. For example, the third line in Eq. (2) indicates that the player decreases the probability of cooperation p_t because selecting D in round t has yielded a satisfactory outcome. In addition, we assume that the player misimplements the action with a small probability ϵ such that the player in fact cooperates with probability $(1 - 2\epsilon)p_t + \epsilon$ in round t . Equations (2) and (3) indicate that the player is satisfied with the current situation if the obtained payoff r_t

is larger than the so-called aspiration level A_t . Otherwise, the player is motivated to flip the action. β controls the sensitivity in the plasticity of p_t . If $\beta = 0$, $s_t = 0$ for any t such that p_t is constant. If $\beta = \infty$, $s_t = 1$ or -1 for any t such that $p_{t+1} = 1$ or 0 .

Unless otherwise stated, we set the initial condition to $p_1 = 0$, i.e., the player defects in round 1. This value of p_1 is the most adverse to mutual cooperation. We will confirm in Secs. 3.1 and 3.4 that our main results are qualitatively the same if we set $p_1 = 1$.

The dynamics of the aspiration level are given by

$$A_{t+1} = (1 - h)A_t + hr_t, \quad (4)$$

where h represents the learning rate. If $h = 0$, A_t is constant, and the model is equivalent to the classical BM model. If $h = 1$, the player compares the current payoff and the payoff obtained in the last round to determine s_t . In our previous work, we showed that mutual cooperation is established among the players only after $t_{\max} = 100$ rounds if β is large and h is small (Masuda & Nakamura, 2011). In the numerical simulations, we set $\beta = 3$, which is large enough to support mutual cooperation if other conditions, such as small h and small ϵ , are met.

We remark that the initial condition A_1 is a key parameter to characterize the player.

2.3 Evolutionary dynamics

We set the number of players in the population to $N = 500$. In a single generation, each player i plays the iterated PD game with $t_{\max} = 200$ against all the other players. We always reset p_t and A_t to p_1 and A_1 when a player starts the iterated PD game with a new opponent. The single generation payoff \bar{r}_i ($\in [S, T]$) is equal to the summation of the payoff obtained by playing against $N - 1$ players, which is divided by $(N - 1)t_{\max}$.

After the single generation payoffs to all the players are determined, we select two players i and j with equal probability for strategy update. We use the Fermi rule (Szabo & Toke, 1998; Traulsen *et al.*, 2006) in which player i adopts j 's A_1 and h values in the next generation with probability $1 / [1 + \exp(\tilde{\beta}(\bar{r}_i - \bar{r}_j))]$, and player j adopts i 's parameter values, otherwise. We set $\tilde{\beta} = 1$. To account for mutation, we assume that after strategy update, A_1 and h of the adopter are displaced by random small values obeying the uniform density on $[-\Delta_{A_1}, \Delta_{A_1}]$

and $[-\Delta_h, \Delta_h]$, respectively. If the displaced h exceeds 1 or is negative, we reset h to 1 or 0, respectively. However, the resetting seldom occurs in our evolutionary simulations.

The phenotype of a player in round t is specified by p_t and A_t . It should be noted that p_t and A_t are not inherited over generations. In other words, the natural selection operates on the capacity to learn (i.e., h) but not on the acquired behavior (i.e., p_t and A_t). Because we let β in Eq. (3) to be relatively large to realize mutual cooperation (Masuda & Nakamura, 2011), p_t is sensitive to the excess payoff relative to A_t in the sense that p_t is close to 0 or 1 unless r_t is close to A_t . Therefore, p_t is similar to the probability of cooperation conditioned on the outcome of the PD game in the previous round. When we use the term learning in the following, we exclusively refer to that induced by h in the iterated PD game. A positive value of h directly raises the plasticity of A_t and indirectly controls that of p_t .

3 Results

3.1 Nash equilibria when without learning

To show that learning is necessary for the emergence of cooperation, we start by analyzing the competition between players that do not learn. With learning rate h equal to zero, the aspiration level is fixed over rounds (i.e., $A_t = A_1$, $t \geq 1$). For the sake of analysis, we set $\beta = \infty$. Then, Eqs. (2) and (3) imply that the player persists in the current action (i.e., $(p_t, p_{t+1}) = (0, 0)$ or $(1, 1)$) if $r_t - A_t \geq 0$ and flips the action (i.e., $(p_t, p_{t+1}) = (0, 1)$ or $(1, 0)$) otherwise. When A_t is fixed, there are five strategies:

- Strategy st1 is defined by $A_t < S$. Except for the action misimplementation, an st1 player always cooperates or always defects, depending on the action in the first round.
- Strategy st2 is defined by $S < A_t < P$. An st2 player does not flip the action unless $r_t = S$.
- Strategy st3 is defined by $P < A_t < R$. An st3 player does not flip the action if mutual cooperation or unilateral defection is realized. It is equivalent to Pavlov, which is a strong competitor in the iterated PD game (Kraines & Kraines, 1989; Nowak & Sigmund, 1993).

- Strategy st4 is defined by $R < A_t < T$. An st4 player flips the action unless $r_t = T$.
- Strategy st5 is defined by $T < A_t$. An st5 player flips the action in every round except when the player misimplements the action.

In Table 1, the average payoff to a nonlearning (i.e., $h = 0$) player (row player) playing against another nonlearning player (column player) is shown for $0 < \epsilon \ll 1$ and $t_{\max} = \infty$. For example, st1 playing against st2 obtains $(R + 3P + 2S)/6$ per round on an average. The results shown in Table 1 are a subset of those obtained in (Nowak *et al.*, 1995) (see Appendix A for details). Table 1 indicates that st3 is a Nash equilibrium when the five strategies are considered. In particular, st3 playing against another st3 realizes mutual cooperation and obtains the largest average payoff per round R . Therefore, a unanimous population composed of st3 players represents a eusocial situation. Mutual cooperation is not realized by any other combination of two players.

Table 1 indicates that st2 is also a Nash equilibrium when $3P > R + 2S$. In addition, although st4 is not a Nash equilibrium, a homogeneous population composed of st4 players is resistant to invasion by st3 in evolutionary situations because st4 gains a larger payoff than st3 does when playing against an st4 opponent.

To test the robustness of the results shown in Table 1, we set $\beta = 3$, $\epsilon = 0.02$, and $h = 0$, and numerically calculate the payoff averaged over t_{\max} generations to different nonlearning players with different fixed aspiration levels A_1 . We also set $R = 4$, $T = 5$, $S = 0$, and $P = 2$ in this and the following numerical simulations. The average payoff to a nonlearner playing against another nonlearner is shown for $p_1 = 0$ and $t_{\max} = 200$ in Fig. 1. The presented values are averages over 100 trials for each pair of A_1 values. The results shown in Fig. 1 are qualitatively the same as those shown in Table 1. We also confirmed that the results hardly change for $(p_1, t_{\max}) = (0, 2000)$, $(1, 200)$, and $(1, 2000)$.

3.2 Possibility of mutual cooperation via learning

If $h > 0$, players different from st3 may adjust A_t until $P < A_t < R$ is satisfied such that they learn to behave as Pavlov. Therefore, learning may play a constructive role in the evolution

of mutual cooperation. In fact, this is not always the case; $\epsilon > 0$ is a necessary condition for mutual cooperation to evolve.

To explain this point, we set $\beta = 3$ and $h = 0.1$, and numerically examine the behavior of a pair of players. Typical time courses of the aspiration level for a pair of learning players over rounds without action misimplementation (i.e., $\epsilon = 0$) are shown in Fig. 2(a). Each of the three pairs with close A_1 values represents a pair of st1 (thick lines), st3 (dotted lines), and st5 players (medium lines), respectively. We used different values of A_1 for each pair for the clarity of the figure; making A_1 equal for two players does not qualitatively change the results. The thick lines in Fig. 2(a) indicate that the two st1 players playing with each other are satisfied with payoff $P = 2$ obtained by mutual defection. Therefore, their aspiration levels converge to $A_t = P$. The results would be the same if we start from a pair of st2 players or a combination of an st1 player and an st2 player. A pair of st3 players begin mutual cooperation from the second round, and their A_t values converge to $R = 4$ (dotted lines). Mutual cooperation is also realized if the two players are initially either st4 or st5, although some rounds are required before the players mutually cooperate (medium lines).

Although two learning players having $A_t < P$ do not end up with mutual cooperation when $\epsilon = 0$, the action misimplementation (i.e., $\epsilon > 0$) can trigger a shift from mutual defection to mutual cooperation. Artificially generated time courses in the presence of action implementation are shown in Fig. 2(b) for expository purposes. Until the intended action is misimplemented ($1 \leq t \leq 29$ in Fig. 2(b)), two players starting with $A_1 < P$ keep mutual defection (thick lines). When A_t has sufficiently approached P , we assume that one player misimplements the action ($t = 30$). Then, the A_t values of both players cross P from below within a couple of rounds such that the players start to behave as Pavlov and mutually cooperate. The possibility of mutual cooperation through this mechanism is sensitive to the value of h . Two players starting with $A_1 < P$ end up with $A_t > P$ owing to the action misimplementation when $(P - S)/(T - P) < (1 - h)^2$ (see Appendix B for derivation). When $T = 5$, $S = 0$, and $P = 2$, this condition yields $0 < h < 1 - \sqrt{2/3} \approx 0.184$ for an arbitrary value of R . Then, the A_t values of the two players converge to R .

The A_t values also converge to R when we start with a pair of st3 players (dotted lines in

Fig. 2(b)) and a pair of st4 or st5 players (medium lines in Fig. 2(b)). This is because mutual cooperation is stable against action misimplementation; if one player turns into D by action misimplementation in round t , both players defect in round $t + 1$ and cooperate in round $t + 2$, if the actions are not misimplemented in rounds $t + 1$ and $t + 2$. This event sequence is likely unless ϵ is large.

3.3 Adaptive dynamics

In the evolutionary numerical simulations that we will describe in Sec. 3.4, we allow the initial aspiration level A_1 and learning rate h to mutate (Sec. 2.3). If the distribution of A_1 and that of h for an evolving population are single peaked and sufficiently localized, we can grasp the evolutionary dynamics for a population by tracking the dynamics of the population averages of A_1 and h , denoted by $\overline{A_1}$ and \overline{h} , respectively. In the extreme case in which all the players share identical values of A_1 and h , the instantaneous dynamics of $\overline{A_1}$ and \overline{h} are captured by adaptive dynamics (Metz *et al.*, 1996; Hofbauer & Sigmund, 1998; Hofbauer & Sigmund, 2003; Doebeli *et al.*, 2004). Adaptive dynamics reveal the possibility for mutants with a slightly deviated parameter value to invade a homogeneous resident population. In this section, we numerically examine two-dimensional adaptive dynamics with respect to $\overline{A_1}$ and \overline{h} to foresee the evolutionary simulations carried out in Sec. 3.4.

In this and the following sections, we set $\beta = 3$, $p_1 = 0$, $\epsilon = 0.02$, and $t_{\max} = 200$ unless otherwise stated. Consider a homogeneous population of players sharing the parameter values $\overline{A_1} = A_1$ and $\overline{h} = h$. A mutant player with aspiration level A'_1 and learning rate h' can invade the population if

$$\pi[s', s] - \pi[s, s] > 0 \quad (5)$$

or

$$\pi[s', s] - \pi[s, s] = 0 \text{ and } \pi[s', s'] - \pi[s, s'] > 0, \quad (6)$$

where $s = (A_1, h)$ and $s' = (A'_1, h')$ are the strategies of the resident and mutant players, respectively, and $\pi[s_1, s_2]$ represents the average payoff of strategy s_1 when playing with strategy s_2 . s is ESS if the converse of Eq. (5) or the converse of Eq. (6) is satisfied. If Eq. (5) or (6)

is satisfied, the homogeneous population comprising strategy s would evolve toward s' . We numerically calculate $\pi[s', s] - \pi[s, s]$, where $s' = (A_1 + 0.2, h)$, $(A_1 - 0.2, h)$, $(A_1, h + 0.02)$, and $(A_1, h - 0.02)$. We confine s' in the neighborhood of s because the amount of mutation for A_1 and h is assumed to be small. Examining $\pi[s', s] - \pi[s, s]$ corresponds to looking at the discretized adaptive dynamics, i.e., the discretized derivative of $\pi[s', s]$ with respect to s' at $s' = s$.

For various values of A_1 and h , $\pi[s', s] - \pi[s, s]$ is shown in Fig. 3. The plotted values are averages over 10^4 runs for any s . In Fig. 3(a), $s' = (A_1 + 0.2, h)$ obtains a larger payoff than $s = (A_1, h)$ in the red region. In this region, s' would invade a homogeneous resident population of s such that $\overline{A_1}$ increases. In contrast, s' obtains a smaller payoff than s does in the blue region. Figure 3(a) indicates that, if learning is prohibited (i.e., $h = 0$), the population starting from $A_1 = 0$, for example, is expected to evolve such that $\overline{A_1}$ increases, but only up to $A_1 \approx P = 2$. Therefore, a population does not evolve from st2 to st3 without learning.

Figure 3(b), which reveals the possibility of invasion by mutant $s' = (A_1 - 0.2, h)$ in the resident population of s , is a sign flipped version of Fig. 3(a) in most parameter regions. Nevertheless, neither the mutants with $A'_1 = A_1 + 0.2$ nor the ones with $A'_1 = A_1 - 0.2$ invade the resident population (i.e., parameter regions colored in blue in both Figs. 3(a) and 3(b)) for $(A_1, h) \approx (2, 0)$ and along a bent line passing through $(A_1, h) \approx (4, 0)$ and $(A_1, h) \approx (4.3, 0.2)$. These regions constitute singular points of the adaptive dynamics and serve as repellers. In other words, $\overline{A_1}$ does not pass through $\approx P = 2$ for $h = 0$ and $\approx R = 4$ for various values of h in adaptive dynamics. The observations for $h = 0$ that the homogeneous population of st2 is not invaded by st3 mutants, that of st3 is not invaded by st2 or st4 mutants, and that of st4 is not invaded by st3 mutants, are consistent with the results obtained in Sec. 3.1.

The possibility of invasion by mutant $s' = (A_1, h + 0.02)$ in the homogeneous population of $s = (A_1, h)$ is shown in Fig. 3(c). The figure suggests that \overline{h} would increase for a population of st1 players (i.e., $A_1 < S = 0$). Learning is preferred to nonlearning when $A_1 < 0$ for the following reason. As shown in Sec. 3.2, when $h > 0$ and $\epsilon > 0$, A_t increases until the players behave as Pavlov to mutually cooperate within a relatively small number of rounds (Fig. 2(b)). In contrast, the players do not establish mutual cooperation when $h = 0$ or $\epsilon = 0$, as shown in

Sec. 3.1 (Fig. 2(a)). Figure 3(c) indicates that \bar{h} increases up to $\bar{h} \approx 0.15$. This value of \bar{h} is consistent with the upper bound of h for mutual cooperation to be possible, which was derived in Sec. 3.2. Based on these results, \bar{h} is expected to initially increase in evolutionary dynamics starting with a population of nonlearning st1 players. We refer to the stage of evolutionary dynamics in which \bar{h} increases as stage 1. The existence of stage 1 is also supported by Fig. 3(d) in which the mutant has $s' = (A_1, h - 0.02)$.

After \bar{h} has increased, Figs. 3(a) and 3(b) imply that \bar{A}_1 increases to cross $P = 2$. When $h > 0$, a larger value of A_1 ($< P$) is beneficial because fewer rounds are required for such players to turn to Pavlov (i.e., $A_t > P$). Once A_1 exceeds P for a majority of players, they earn a large average payoff $\approx R$ through mutual cooperation. We refer to the transition for learning players from a small A_1 corresponding to st1 or st2 to a large A_1 corresponding to st3 as stage 2. Figures 3(a) and 3(b) indicate that the difference between $\pi[s', s]$ and $\pi[s, s]$ when $A_1, A_1 + 0.2 < P, h > 0$ is small, presumably because s and s' are only slightly different in terms of the number of transient rounds before the entrance to $A_1 > P$. Therefore, we expect that stage 2 occurs slowly in evolutionary dynamics.

Although it is a minor phenomenon as compared to stages 1 and 2, a smaller h is more beneficial on the boundary between st2 and st3 (i.e., $A_t \approx P = 2$), as shown in Figs. 3(c) and 3(d). For expository purposes, time courses of the iterated PD game between an st2 player and an st3 player are shown in Fig. 3(e). As shown by the solid lines, the initial st3 player flips to st2 before establishing mutual cooperation if $h > 0$. In fact, a nonlearning st3 player (i.e., $h = 0$) realizes mutual cooperation with a learning st2 player in earlier rounds (dotted lines) than a learning st3 player does (solid lines). Therefore, in evolutionary dynamics, \bar{h} in the vicinity of $A_t \approx P$ is expected to decrease. We refer to this transition as stage 3. It should be noted that stage 3 occurs in a narrow range of A_1 (i.e., $A_1 < P$ and $A_1 + 0.2 > P$ in Fig. 3(a)).

Through stages 1, 2, and 3, evolution from a defective population of nonlearning st1 players to a cooperative population of st3 players is logically possible. In contrast, the emergence of mutual cooperation is hampered if learning is prohibited.

After stage 3, \bar{A}_1 would not evolve beyond R ; $A_1 \approx R$ is a line of repellers in adaptive dynamics, as already explained in Figs. 3(a) and 3(b). When $P < A_1 < R$ (i.e., $2 < A_1 < 4$),

the mutant's payoff is indistinguishable from the resident's payoff unless h is large (Figs. 3(a)–(d)). Therefore, $\overline{A_1}$ and \overline{h} would perform approximately unbiased diffusion. This implies that \overline{h} that has decreased via stage 3 may increase again.

When at least one of the two players is st4 or st5, a player with a larger A_1 is more advantageous than the opponent with a smaller A_1 . This is because the former exploits the latter in early rounds. Nevertheless, these players do not obtain the average payoff as large as that for a pair of st3 players, which would start to mutually cooperate from the second round. Therefore, st3 is stable against invasion by st4 and vice versa.

We predict that the learning rate would not eventually decrease to the small value in evolutionary simulations. In other words, the disadvantage of learning is too small to be evolutionarily relevant unless the cost of learning is explicitly incorporated.

To assess the robustness of the results obtained from the adaptive dynamics, we reproduced Figs. 3(a)–(d) with $\epsilon = 0.05$ and $\epsilon = 0.1$. The results for $\epsilon = 0.05$ are qualitatively the same as those for $\epsilon = 0.02$ (results not shown). The results for $\epsilon = 0.1$ are different in some aspects from those for $\epsilon = 0.02$ (Fig. 4). Most notably, when $\epsilon = 0.1$, st3 is no longer stable against invasion by st2 even without learning (i.e., $h = 0$). Therefore, mutual cooperation would not be stable in evolutionary dynamics. In Figs. 5(a) and 5(b), $\pi[s', s] - \pi[s, s]$ is shown for $(s, s') = ((1.9, 0), (2.1, 0))$ and $(s, s') = ((2.1, 0), (1.9, 0))$, respectively, for a variety of values of t_{\max} and ϵ . Figure 5(a) indicates that an st3 mutant does not invade the population of st2 residents for all the examined values of t_{\max} and ϵ . Figure 5(b) indicates that a population of st3 residents is resistant to invasion by st2 mutants when t_{\max} is large and $\epsilon (> 0)$ is small. Nevertheless, st3 is stable for various values of t_{\max} and ϵ . Because stage 2 is hampered when $\epsilon = 0$, ϵ must take an intermediate value for the learning-mediated mutual cooperation to emerge.

3.4 Evolutionary simulations

The results in Sec. 3.3 predict the presence of a learning mediated evolutionary route from a noncooperative population composed of st1 players to a cooperative population composed of st3 players. In this section, we carry out direct numerical simulations of the evolutionary

dynamics using a population composed of $N = 500$ players. We initially set $h = 0$ and select A_1 for each player independently from the uniform density on $[S - 1, S]$. Therefore, all the players are initially nonlearning st1. Refer to Sec. 2.3 for details of the numerical setup.

The evolution of \bar{h} , the total amount of plasticity experienced in a generation, defined by $\sum_{t=1}^{t_{\max}-1} |A_{t+1} - A_t|$, \bar{r} , and the fraction of mutual cooperation for an example run with $\Delta_{A_1} = 0.05$ and $\Delta_h = 0.01$ are shown in Fig. 6(a). The average learning rate \bar{h} and the total amount of plasticity rapidly increase until the $\approx 3.9 \times 10^4$ th round. The payoff and the fraction of mutual cooperation also increase during this period because st1 players learn to behave as Pavlov when $h > 0$. This period corresponds to stage 1 described in Sec. 3.3. Then, the fraction of mutual cooperation and the total amount of plasticity gradually increase until the $\approx 3.5 \times 10^5$ th round, corresponding to stage 2. In the $\approx 3.5 \times 10^5$ th round, an st3 mutant emerges in the population mostly composed of st2 players and gains a larger payoff than st2 residents do. Then, st3 players rapidly replace st2 players in the population such that \bar{r}/R and the fraction of mutual cooperation suddenly increase (Fig. 6(a)). This is because stable mutual cooperation between st3 players emerges in an early round, whereas that between st2 players emerges after $\approx 2/\epsilon$ rounds. The learning rate decreases almost at the same time, corresponding to stage 3. The time courses of the fractions of st1, st2, st3, st4, and st5 players corresponding to the run shown in Fig. 6(a) are shown in Fig. 6(b). For example, the fraction of the st1 player is defined by the fraction of players having $A_1 < S$ and any value of h . Figure 6(b) indicates that the population initially composed of st1 players evolves to that of st3 players. The trajectory of \bar{A}_1 and \bar{h} corresponding to the same run is shown in Fig. 6(c). Figure 6(c) is consistent with the scenario of the evolution of cooperation described in Sec. 3.3. The population evolves from no cooperation to mutual cooperation via the three stages involving learning. After stage 3, \bar{A}_1 and \bar{h} diffuse without a recognizable bias, which is also consistent with the results obtained in Sec. 3.3 (white regions in Figs. 3(a)–(d)). However, it should be noted that the total amount of plasticity remains small after stage 3.

To examine the robustness of the results, we carry out five runs of numerical simulations for each of the different parameter sets; we could not carry out more extensive numerical simulations because of the computational cost. We measure two quantities in each run. The

first quantity is the number of generations necessary for \bar{h} to exceed 0.1 for the first time. We call this number the end of stage 1. The second quantity is the number of generations necessary for \bar{A}_1 to exceed P for the first time. We call this number the end of stage 2. The ends of stages 1 and 2 with $\epsilon = 0.02$, $\Delta_{A_1} = 0.05$, and $\Delta_h = 0.01$ are equal to 257 ± 28 ($\times 10^2$) and 3763 ± 1514 ($\times 10^2$), respectively, where the mean \pm standard deviation on the basis of the five runs are indicated. Those with $\epsilon = 0.05$, $\Delta_{A_1} = 0.05$, and $\Delta_h = 0.01$ are equal to 299 ± 55 ($\times 10^2$) and 6951 ± 4231 ($\times 10^2$). Those with $\epsilon = 0.01$, $\Delta_{A_1} = 0.05$, and $\Delta_h = 0.01$ are equal to 265 ± 31 ($\times 10^2$) and 3780 ± 1547 ($\times 10^2$). Those with $\epsilon = 0.02$, $\Delta_{A_1} = 0.02$, and $\Delta_h = 0.01$ are equal to 239 ± 18 ($\times 10^2$) and 9151 ± 6836 ($\times 10^2$). For this parameter set, one out of the five runs did not reach the end of stage 2 within 2×10^6 generations, such that the statistics are based on the other four runs. Those with $\epsilon = 0.02$, $\Delta_{A_1} = 0.05$, and $\Delta_h = 0.005$ are equal to 613 ± 119 ($\times 10^2$) and 3504 ± 999 ($\times 10^2$). Mutual cooperation evolves via learning (i.e., finite value of the end of stage 2 up to our numerical efforts) in most cases. When $\epsilon = 0.05$, evolution to mutual cooperation is slower than when $\epsilon = 0.02$. This may be because learning players having different values of A_1 turn into Pavlov (i.e., $A_t > P$) within a small number of rounds when ϵ is relatively large. Then, the payoff to different learning players would differ relatively little to weaken the selection pressure.

We perform another robustness test. For the original parameter values $\epsilon = 0.02$, $\Delta_{A_1} = 0.05$, and $\Delta_h = 0.01$, the trajectory of \bar{A}_1 and \bar{h} obtained from a single run with $p_1 = 1$ is shown in Fig. 7. The results are qualitatively the same as those for $p_1 = 0$ (Fig. 6(c)) although establishment of cooperation takes a considerably larger number of generations when $p_1 = 1$ than when $p_1 = 0$.

3.5 Baldwin effect

If we assume an explicit cost of learning, the learning rate decreases after mutual cooperation is reached. An example time course of \bar{A}_1 and \bar{h} when a linear cost $-ch$ is added to the single generation payoff to each player (Suzuki & Arita 2004; see Ancel 1999, 2000 for a different implementation of the explicit learning cost), where $c = 1$, is shown in Fig. 8(a). The final value of \bar{h} is smaller than that in the case without the learning cost (Fig. 6(c)). The result

shown in Fig. 6(d) is an example of the standard Baldwin effect in which the learning rate initially increases and then decreases (Ancel, 2000; Dopazo *et al.*, 2001; Borenstein *et al.*, 2006; Paenke *et al.*, 2007; Paenke *et al.*, 2009).

We examine the robustness of the observed Baldwin effect against the variation of c . An example time course of $\overline{A_1}$ and \overline{h} when $c = 10$ is shown in Fig. 8(b). As compared to when $c = 1$ (Fig. 8(a)), \overline{h} is smaller throughout the evolution, and $\overline{A_1}$ increases more slowly. Nevertheless, the population mostly consists of st3 players in the end. We carried out five runs for various values of c . The largest \overline{h} value in 10^6 generations is shown in Fig. 8(c) for each run. The largest \overline{h} value decreases with c because learning is costly for a large value of c . The final value of \overline{h} , calculated as the average over the last 10^4 generations, is shown in Fig. 8(d). The final value of \overline{h} is considerably smaller than the largest value (Fig. 8(c)) for each c , indicative of stage 3 of the Baldwin effect. The final value of $\overline{A_1}$, calculated as the average over the last 10^4 generations, is plotted against c in Fig. 8(e). If this value is larger than $P = 2$ and smaller than $R = 4$, we expect that the final population is mostly composed of st3 players and that the Baldwin effect is operative. Figure 8(e) suggests that the Baldwin effect occurs in the five runs when $c < 15$. When $c \geq 15$, stage 1, i.e., the initial increase in \overline{h} , is often too small in magnitude such that stage 2 does not sometimes occur. We conclude that the Baldwin effect occurs for a wide range of c .

4 Discussion

We have shown that reinforcement learning promotes the evolution of mutual cooperation in a population of players involved in the iterated PD game. Cooperation evolves under some conditions such as $3P > R + 2S$, positive but not too large values of ϵ , and t_{\max} that is not too small. The present study is motivated by previous investigations of the Baldwin effect. Our results provide an example of the Baldwin effect in the form of a computational model of social behavior.

To understand the behavior of our model analytically, writing down the Fokker–Planck equation for the joint density of A_1 and h may be useful. Starting from the singular density at

a small value of A_1 and $h = 0$, we may be able to solve the Fokker–Planck equation numerically to track the evolution of the joint density to find the Baldwin effect. Alternatively, discretizing A_1 and h and then formulating a Markov chain on the discretized states may also be useful. Nevertheless, we refrained from such analyses because we consider that they eventually necessitates some numerical simulations and would not sufficiently advance the understanding of our numerical results.

The concept of the Baldwin effect is diverse (Simpson, 1953; Downes, 2003; Turney *et al.*, 1996). However, arguably, the most accepted variant of the Baldwin effect is formulated as a two-stage mechanism (Simpson, 1953; Godfrey-Smith, 2003; Crispo, 2007; Turney *et al.*, 1996). In stage 1, plasticity increases because plastic individuals are better at finding the optimal behavior than nonplastic individuals. In stage 2, mutation makes the optimal behavior innate and decreases the plasticity of individuals. Mutants that play optimally from the outset of their life without plasticity and resident individuals that acquire the optimal behavior through plasticity are eventually equally efficient. Nevertheless, because of the cost of learning, the mutants overwhelm the residents via natural selection. Stage 2 is often called genetic assimilation.

Stage 1 in our model corresponds to stage 1 of the standard Baldwin effect outlined above. In stage 2 in our model, A_1 increases such that the optimal behavior (i.e., mutual cooperation by turning into st3) becomes innate. Nevertheless, after stage 3 in our model in which the learning rate rapidly decreases, the learning rate starts to perform a random walk because the learning cost is marginal in our model (Fig. 6(c)). Therefore, the behavior of our model in stages 2, 3, and onward does not qualify as stage 2 of the standard Baldwin effect in which the learning rate decreases. With a modified model with an explicit learning cost, we showed that the learning rate decreases after stage 3 (Fig. 6(d)). In this case, our model naturally fits the framework of the Baldwin effect.

In a previous computational model of the Baldwin effect, learning rates remain large when the optimal behavior dynamically changes owing to environmental fluctuations (Ancel, 1999). In our model without an explicit cost of learning, the learning rate remains large for a different reason. In our model, the optimal parameter set (i.e., A_1 and h) does not fluctuate after sufficient generations. Instead, approximate optimality is realized for various parameter sets,

i.e., any $P < A_1 < R$ and $h \geq 0$. Therefore, the learning rate performs a random walk to occasionally visit large values (Fig. 6(c)).

Godfrey-Smith points out three alternative reasons why stage 1 cannot be skipped in the two-stage mechanism of the Baldwin effect (Godfrey-Smith, 2003). First, learning may provide a breathing space by which a population can survive long enough to transit to stage 2. This reason is irrelevant to our model because our model is not concerned with the survival of the population. The population size is fixed in our model such that the population always survives. Second, the preferred state may be accessible for learners but not for nonlearners. Although not explicitly stated in Godfrey-Smith (2003), this mechanism seems to be relevant to cases in which the fitness landscape does not depend on the configuration of the population. In our case, however, the fitness landscape depends on the fractions of the different types of players because the payoff to a player is affected by the strategies of the other players. Third, evolution may change the “social ecology” of the population such that learners are more advantageous than nonlearners, a phenomenon called niche construction in a broad sense. The social ecology implies a fitness landscape that depends on the configuration of the population. In our model, the social ecology evolves via learning of players. This third mechanism seems to be relevant to our model. Suppose a hypothetical population comprising $st1$ nonlearners except two $st1$ learners. For a focal $st1$ learner, the social ecology is such that there is one $st1$ learner and $N - 2$ $st1$ nonlearners. If $\epsilon > 0$, the focal $st1$ learner is likely to gain a payoff that is larger than an $st1$ nonlearner because the focal player learns to mutually cooperate with the other $st1$ learner, whereas an $st1$ nonlearner does not. The focal $st1$ learner would not overwhelm $st1$ nonlearners if the other $st1$ learner is absent in the social ecology.

The main purpose of this study is to provide an evolutionary model of concrete social behavior in which learning plays a constructive role. We are not the first to achieve this end. Suzuki and Arita observed the Baldwin effect in the iterated PD game using different learning models (Suzuki & Arita, 2004). In their model, the learning rate is assumed to be binary, and the player’s strategy is specified by a look-up table that associates the action to take (i.e., C or D) with the actions of the previous two rounds of the two players. The entries of the look-up table dynamically change when the plasticity is in operation. They also considered the effects

of meta-learning in which the player adapts how to update each entry of the look-up table. The main contribution of the present work relative to theirs is to provide a much simpler model in terms of the number of plastic parameters. In contrast, the learning rates and the range of parameters are continuous in our model, whereas they are mostly binary in their model. Our model may be amenable to real animals and facilitates a mechanistic understanding of evolutionary dynamics by the numerically calculated adaptive dynamics. Apart from the fixed parameters common to all the individuals, our players only have two parameters that are plastic within a generation, p_t and A_t , and two parameters inherited across generations, A_1 and h . The results obtained from the adaptive dynamics predict those of direct evolutionary numerical simulations and provide an intuitive reason why learning promotes the emergence of mutual cooperation. In particular, we showed the necessity of the evolution of learning ability for cooperation by explicitly comparing the cases with and without learning. The combination of adaptive dynamics and evolutionary simulations may also be useful for analyzing the Baldwin effect in different models.

Acknowledgements

We thank Reiji Suzuki and Kohei Tamura for the valuable discussions. N.M. acknowledges the support provided through Grants-in-Aid for Scientific Research (Nos. 20760258 and 23681033, and Innovative Areas “Systems Molecular Ethology” (No. 20115009)) from MEXT, Japan.

Appendix A: Payoff to nonlearning players

Nowak *et al.* (1995) analyzed iterated matrix games between a pair of players that select an action (i.e., C or D) in response to the actions of the two players in the previous round. There are four combinations of the actions of the two players in the previous round, i.e., (C, C), (C, D), (D, C), and (D, D). Because a player assigns C or D to each of these possible outcomes in the previous round, there are 16 strategies S_i ($0 \leq i \leq 15$). In fact, st1, st2, st3, st4, and st5 in the present study are equivalent to S_{12} , S_8 , S_9 , S_1 , and S_3 in (Nowak *et al.*, 1995), respectively.

By calculating the steady state of the Markov chain with four states R , T , S , and P , Nowak *et al.* (1995) calculated the average payoff to focal player S_i playing against the opponent S_j ($0 \leq i, j \leq 15$) under a small probability of error in action implementation. Their assumption for the action misimplementation is slightly different from ours. We assumed that ϵ is the probability that each player independently misimplements the action, whereas only one of the two players may misimplement the action in a round in their model. Nevertheless, our model is equivalent to theirs in the limit $\epsilon \rightarrow 0$ if we set $\epsilon' = 2\epsilon(1 - \epsilon)$, where ϵ' is the probability of action misimplementation in the sense of Nowak *et al.* (1995). Therefore, our results shown in Table 1 are a corollary of their results.

Appendix B: Upper bound of h for st2 players to turn into Pavlov

Given $\beta = \infty$, $0 < \epsilon \ll 1$, $h > 0$, and $A_1 < P$, A_t of the two players, denoted by X and Y, are sufficiently close to P when one player, which we assume to be Y without loss of generality, misimplements the action to select C for the first time in round $t \propto 2/\epsilon$. Without any further action misimplementation, X keeps D and Y flips to D in round $t+1$ because $A_{t+1}^{(Y)} < P < A_{t+1}^{(X)}$. In round $t+2$, X flips to C and Y keeps D. Therefore, we obtain $A_{t+3}^{(X)} = hS + (1 - h)A_{t+2}^{(X)}$, $A_{t+2}^{(X)} = hP + (1 - h)A_{t+1}^{(X)}$, $A_{t+1}^{(X)} = hT + (1 - h)A_t^{(X)}$, and $A_t^{(X)} \approx P$. Combining the four equations, we obtain

$$A_{t+3}^{(X)} = (T - P)h^3 - 2(T - P)h^2 + (T + S - 2P)h + P. \quad (7)$$

Using $T > P > S$ and $0 < h \leq 1$, we obtain the condition for X to become Pavlov in round $t+3$ as $A_{t+3}^{(X)} > P$, i.e.,

$$(1 - h)^2 > \frac{P - S}{T - P}. \quad (8)$$

The condition for Y to become Pavlov in round $t+3$ is given by $A_{t+3}^{(Y)} > P$, i.e.,

$$(1 - h)^2 < \frac{T - P}{P - S}. \quad (9)$$

Equation (8) implies Eq. (9) because $(P - S)/(T - P) > 0$ and $h > 0$. Therefore, the two players become Pavlov in round $t + 3$ if Eq. (8) holds true.

We assume that Eq. (8) is violated. If Eq. (9) is also violated, we obtain $A_{t'}^{(X)}, A_{t'}^{(Y)} \leq P$ ($t' \geq t + 3$) such that the two players mutually defect until the occurrence of another action misimplementation. If Eq. (9) is satisfied, the two players mutually defect in round $t + 3$. Because $r_{t+2}^{(Y)} = r_t^{(X)} = T$, $r_{t+3}^{(Y)} = r_{t+1}^{(X)} = P$, and $A_{t+2}^{(Y)} < P$, we obtain $P < A_{t+4}^{(Y)} < A_{t+2}^{(X)}$. Because $r_{t+2}^{(X)} = r_t^{(Y)} = S$, $r_{t+3}^{(X)} = r_{t+1}^{(Y)} = P$, and $A_{t+2}^{(X)} > P$, we obtain $A_{t+2}^{(Y)} < A_{t+4}^{(X)} < P$. These two inequalities indicate that X and Y behave as st3 and st2 in round $t + 5$, respectively. By repeating the same procedure with X and Y swapped, we obtain

$$A_{t+2}^{(Y)} < A_{t+6}^{(Y)} < P < A_{t+6}^{(X)} < A_{t+2}^{(X)}. \quad (10)$$

Therefore, we obtain

$$A_{t+2}^{(Y)} < A_{t+2+4i}^{(Y)} < P < A_{t+2+4i}^{(X)} < A_{t+2}^{(X)} \quad (i \geq 1) \quad (11)$$

by induction. Equation (11) implies that the two players do not realize mutual cooperation if Eq. (8) is violated.

Therefore, an upper bound of h for a pair of st2 players to turn into Pavlov is given by solving Eq. (8) with equality.

References

Ancel, L. W. 1999. A quantitative model of the Simpson–Baldwin effect. *J. Theor. Biol.* 196, 197–209.

Ancel, L. W. 2000. Undermining the Baldwin expediting effect: does phenotypic plasticity accelerate evolution? *Theor. Popul. Biol.* 58, 307–319.

Axelrod, R. 1984. *The Evolution of Cooperation*. Basic Books, NY.

Badyaev, A. V. 2009. Evolutionary significance of phenotypic accommodation in novel environments: an empirical test of the Baldwin effect. *Philos. Trans. R. Soc. B* 364, 1125–1141.

Borenstein, E., Meilijson, I. & Ruppin, E. 2006. The effect of phenotypic plasticity on evolution in multipeaked fitness landscapes. *J. Evol. Biol.* 19, 1555–1570.

Camerer, C. F. 2003. *Behavioral Game Theory*. Princeton University Press, NJ.

Crispo, E. 2007. The Baldwin effect and genetic assimilation: revisiting two mechanisms of evolutionary change mediated by phenotypic plasticity. *Evolution* 61, 2469–2479.

Doebeli, M., Hauert, C. & Killingback, T. 2004. The evolutionary origin of cooperators and defectors. *Science* 306, 859–862.

Dopazo, H., Gordon, M. B., Perazzo, R. & Risau-Gusman, S. 2001. A model for the interaction of learning and evolution. *Bull. Math. Biol.* 63, 117–134.

Downes, S. M. 2003. *Baldwin effects and the expansion of the explanatory repertoire in evolutionary biology*. In: *Evolution and learning — The Baldwin effect reconsidered* pp. 33–51. MIT Press, Cambridge, UK.

Fudenberg, D. & Levine, D. K. 1998. *The Theory of Learning in Games*. MIT Press, Cambridge, UK.

Godfrey-Smith, P. 2003. *Between Baldwin skepticism and Baldwin boosterism*. In: *Evolution and learning — The Baldwin effect reconsidered* pp. 53–67. MIT Press, Cambridge, UK.

Hinton, G. E. & Nowlan, S. J. 1987. How learning can guide evolution. *Complex Syst.* 1, 495–502.

Hofbauer, J. & Sigmund, K. 1998. *Evolutionary Games and Population Dynamics*. Cambridge University Press, Cambridge, UK.

Hofbauer, J. & Sigmund, K. 2003. Evolutionary game dynamics. *Bull. Am. Math. Soc.* 40, 479–519.

Karandikar, R., Mookherjee, D., Ray, D. & Vega-Redondo, F. 1998. Evolving aspirations and cooperation. *J. Econ. Theory* 80, 292–331.

Kraines, D. & Kraines, V. 1989. Pavlov and the prisoner's dilemma. *Theory Decis.* 26, 47–79.

Lindgren, K. 1991. *Evolutionary phenomena in simple dynamics.* In: *Proceedings of Artificial Life II* pp. 295–312.

Macy, M. 1996. Natural selection and social learning in prisoner's dilemma: co-adaptation with genetic algorithms and artificial neural networks. *Sociol. Methods Res.* 25, 103–137.

Macy, M. W. & Flache, A. 2002. Learning dynamics in social dilemmas. *Proc. Natl. Acad. Sci. USA* 99, 7229–7236.

Masuda, N. & Nakamura, M. 2011. Numerical analysis of a reinforcement learning model with the dynamic aspiration level in the iterated prisoner's dilemma. *J. Theor. Biol.* 278, 55–62.

Masuda, N. & Ohtsuki, H. 2009. A theoretical analysis of temporal difference learning in the iterated prisoner's dilemma game. *Bull. Math. Biol.* 71, 1818–1850.

Maynard Smith, J. 1987. Natural selection: when learning guides evolution. *Nature* 329, 761–762.

Metz, J. A. J., Geritz, S. A. H., Meszena, G., Jacobs, F. J. A. & van Heerwarden, J. S. 1996. *Adaptive dynamics: a geometrical study of the consequences of nearly faithful reproduction.* In: *Stochastic and spatial structures of dynamical systems (S. J. van Strien and S. M. Verduyn Lunel eds.)* pp. 188–231. North Holland, Amsterdam, The Netherlands.

Nowak, M. & Sigmund, K. 1993. A strategy of win-stay, lose-shift that outperforms tit-for-tat in the prisoner's dilemma game. *Nature* 364, 56–58.

Nowak, M. A. & Sigmund, K. 1992. Tit for tat in heterogeneous populations. *Nature* 355, 250–253.

Nowak, M. A., Sigmund, K. & El-Sedy, E. 1995. Automata, repeated games and noise. *J. Math. Biol.* 33, 703–722.

Paenke, I., Kawecki, T. J. & Sendhoff, B. 2009. The influence of learning on evolution: a mathematical framework. *Artif. Life* 15, 227–245.

Paenke, I., Sendhoff, B. & Kawecki, T. J. 2007. Influence of plasticity and learning on evolution under directional selection. *Am. Nat.* 170, E47–E58.

Posch, M., Pichler, A. & Sigmund, K. 1999. The efficiency of adapting aspiration levels. *Proc. R. Soc. Lond. B* 266, 1427–1435.

Sandholm, T. W. & Crites, R. H. 1996. Multiagent reinforcement learning in the iterated prisoner’s dilemma. *Biosystems* 37, 147–166.

Simpson, G. G. 1953. The Baldwin effect. *Evolution* 7, 110–117.

Suzuki, R. & Arita, T. 2004. Interactions between learning and evolution: outstanding strategy generated by the Baldwin effect. *Biosystems* 77, 57–71.

Szabo, G. & Toke, C. 1998. Evolutionary prisoner’s dilemma game on a square lattice. *Phys. Rev. E* 58, 69–73.

Taiji, M. & Ikegami, T. 1999. Dynamics of internal models in game players. *Physica D* 134, 253–266.

Traulsen, A., Nowak, M. A. & Pacheco., J. M. 2006. Stochastic dynamics of invasion and fixation. *Phys. Rev. E* 74, 011909.

Trivers, R. L. 1971. The evolution of reciprocal altruism. *Q. Rev. Biol.* 46, 35–57.

Turney, P., Whitley, D. & Anderson, R. W. 1996. Evolution, learning, and instinct: 100 years of the Baldwin effect. *Evol. Comput.* 4, 4–8.

Waddington, C. H. 1942. Canalization of development and the inheritance of acquired characters. *Nature* 150, 563–565.

Wang, S., Szalay, M. S., Zhang, C. & Csermely, P. 2008. Learning and innovative elements of strategy adoption rules expand cooperative network topologies. *PLoS One*, 3, e1997.

Weber, B. H. & Depew, D. J. eds. 2003. *Evolution and learning — The Baldwin effect reconsidered*. MIT Press, Cambridge, UK.

Yeh, P. J. & Price, T. D. 2004. Adaptive phenotypic plasticity and the successful colonization of a novel environment. Am. Nat. 164, 531–542.




Figure 1: Average payoff to a nonlearning player (row player) playing against an opponent nonlearning player (column player). We set $\beta = 3$, $p_1 = 0$, $\epsilon = 0.02$, $h = 0$, $t_{\max} = 200$, $R = 4$, $T = 5$, $S = 0$, and $P = 2$.




Figure 2: Behavior of a pair of learning players. We set $\beta = 3$, $p_1 = 0$, and $h = 0.1$. (a) Example time courses of the aspiration level for a pair of players when $\epsilon = 0$. The horizontal lines represent $A_t = P = 2$ and $A_t = R = 4$. We set A_1 for the two players to -1 and -0.5 (thick lines), 2.5 and 3 (dotted lines), and 5.5 and 6 (medium lines). (b) Example time courses of the aspiration level when $\epsilon = 0.02$. We set A_1 as in (a). For each pair, one of the two players is assumed to misimplement the action to cooperate in round 30 (thick lines), to defect in round 30 (dotted lines), or to cooperate in round 7 (medium lines).




Figure 3: Discretized adaptive dynamics when $\epsilon = 0.02$. Plotted is $\pi[s', s] - \pi[s, s]$, where $s = (A_1, h)$ and (a) $s' = (A_1 + 0.2, h)$, (b) $s' = (A_1 - 0.2, h)$, (c) $s' = (A_1, h + 0.02)$, and (d) $s' = (A_1, h - 0.02)$. (e) Example time courses of a pair of players having $(A_1, h) = (1.7, 0.1)$ and $(A_1, h) = (2.2, 0.1)$ (solid lines), and $(A_1, h) = (1.7, 0.1)$ and $(A_1, h) = (2.2, 0)$ (dotted lines). The horizontal line indicates $A_t = P = 2$. We set $\beta = 3$, $p_1 = 0$, and $t_{\max} = 200$.




Figure 4: Discretized adaptive dynamics when $\epsilon = 0.1$. Plotted is $\pi[s', s] - \pi[s, s]$, where $s = (A_1, h)$ and (a) $s' = (A_1 + 0.2, h)$, (b) $s' = (A_1 - 0.2, h)$, (c) $s' = (A_1, h + 0.02)$, and (d) $s' = (A_1, h - 0.02)$.




Figure 5: Effects of t_{\max} and ϵ on the adaptive dynamics when $A_1 \approx P$ and $h = 0$. We set $\beta = 3$ and $p_1 = 0$. (a) $\pi[(2.1, 0), (1.9, 0)] - \pi[(1.9, 0), (1.9, 0)]$. (b) $\pi[(1.9, 0), (2.1, 0)] - \pi[(2.1, 0), (2.1, 0)]$.




Figure 6: Evolutionary dynamics in a population composed of learning players. We set $\beta = 3$, $p_1 = 0$, $\epsilon = 0.02$, $t_{\max} = 200$, $\Delta_{A_1} = 0.05$, and $\Delta_h = 0.01$. (a) Time course of \bar{h} , $\sum_{t=1}^{t_{\max}-1} |A_{t+1} - A_t| / 20$, \bar{r}/R , and the fraction of mutual cooperation. (b) Time course of the fraction of initially st1, st2, st3, st4, and st5 players in the run shown in (a). (c) Sample trajectory of the population averages \bar{A}_1 and \bar{h} in the run shown in (a).




Figure 7: Sample trajectory of $\overline{A_1}$ and \overline{h} when $\beta = 3$, $p_1 = 1$, $\epsilon = 0.02$, $t_{\max} = 200$, $\Delta_{A_1} = 0.05$, and $\Delta_h = 0.01$.




Figure 8: (a, b) Sample trajectory of \overline{A}_1 and \overline{h} under a linear cost of learning. We set (a) $c = 1$ and (b) $c = 10$. (c) Largest \overline{h} in each of the five runs of 10^6 generations for various values of c . A cross corresponds to a single run. (d) Average of \overline{h} over the last 10^4 generations in each run. (e) Average of \overline{A}_1 over the last 10^4 generations in each run.

Table 1: Average payoff to a nonlearning player (row player) playing against an opponent nonlearning player (column player). We set $\beta = \infty$, $0 < \epsilon \ll 1$, and $t_{\max} = \infty$.

	st1	st2	st3	st4	st5
st1	$\frac{R+T+S+P}{4}$	$\frac{R+2S+3P}{6}$	$\frac{R+T+S+P}{4}$	$\frac{T+2S+P}{4}$	$\frac{R+T+S+P}{4}$
st2	$\frac{R+2T+3P}{6}$	P	$\frac{R+2T+2P}{5}$	$\frac{T+P}{2}$	$\frac{T+P}{2}$
st3	$\frac{R+T+S+P}{4}$	$\frac{R+2S+2P}{5}$	R	$\frac{R+S+P}{3}$	$\frac{R+T+S+P}{4}$
st4	$\frac{2T+S+P}{4}$	$\frac{S+P}{2}$	$\frac{R+T+P}{3}$	$\frac{R+P}{2}$	$\frac{R+P}{2}$
st5	$\frac{R+T+S+P}{4}$	$\frac{S+P}{2}$	$\frac{R+T+S+P}{4}$	$\frac{R+P}{2}$	$\frac{R+T+S+P}{4}$